arXiv:1703.10956v1 [cs.CV] 31 Mar 2017

InverseFaceNet: Deep Single-Shot Inverse Face Rendering From A Single Image

Hyeongwoo Kim[1] Michael Zollhöfer[1] Ayush Tewari[1]
Justus Thies[2] Christian Richardt[3] Christian Theobalt[1]
[1] Max-Planck-Institute for Informatics [2] University of Erlangen-Nuremberg [3] University of Bath

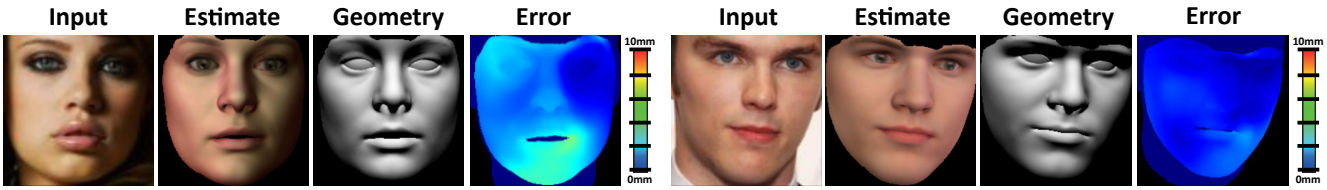

Figure 1. Our single-shot deep inverse face renderer *InverseFaceNet* obtains a high-quality geometry, reflectance and illumination estimate from just a single input image. We jointly recover the face pose, shape, expression, reflectance and incident scene illumination. *From left to right:* input photo, our estimated face model, its geometry, and the pointwise Euclidean error compared to Garrido et al. [14].

Abstract

We introduce InverseFaceNet, a deep convolutional inverse rendering framework for faces that jointly estimates facial pose, shape, expression, reflectance and illumination from a single input image in a single shot. By estimating all these parameters from just a single image, advanced editing possibilities on a single face image, such as appearance editing and relighting, become feasible. Previous learning-based face reconstruction approaches do not jointly recover all dimensions, or are severely limited in terms of visual quality. In contrast, we propose to recover high-quality facial pose, shape, expression, reflectance and illumination using a deep neural network that is trained using a large, synthetically created dataset. Our approach builds on a novel loss function that measures model-space similarity directly in parameter space and significantly improves reconstruction accuracy. In addition, we propose an analysis-by-synthesis 'breeding' approach which iteratively updates the synthetic training corpus based on the distribution of real-world images, and we demonstrate that this strategy outperforms completely synthetically trained networks. Finally, we show high-quality reconstructions and compare our approach to several state-of-the-art approaches.

1. Introduction

Inverse rendering aims to reconstruct scene properties such as geometry, reflectance and illumination from image data. This reconstruction is fundamentally challenging, as it inevitably requires inverting the complex real-world image formation process. It is also an ill-posed problem as certain effects, such as low-frequency reflectance and illumination, can be indistinguishable [34]. Inverse rendering, for example, enables relighting of faces by modifying the scene illumination and keeping the face reflectance and geometry fixed.

Recently, optimization-based approaches for inverse face rendering were introduced with convincing results [2, 14, 21, 26, 45]. One of the key ingredients that enables to disentangle pose, geometry (both related to shape and facial expression), reflectance and illumination are specific priors that constrain parameters to plausible values and distributions. Formulating such priors accurately for real faces is difficult, as they are unknown a priori. The priors could be learned by applying inverse rendering to a large dataset of real face images, but this is highly challenging without having the priors a priori.

We take a different approach to avoid this chicken-and-egg problem. Instead of formulating explicit priors, we directly learn inverse face rendering with a deep neural network that implicitly learns priors based on the training corpus. As annotated training data is hard to come by, we adopt an analysis-by-synthesis approach that starts from a large dataset of synthetic face images with known model parameters (geometry, reflectance and illumination). This is similar to Richardson et al.'s approaches [35, 36], but the parameter distribution of this training data does not match that of real-world faces and environments. As a result, the learned implicit priors are rather weak and do not generalize well to

in-the-wild images.

In contrast to previous approaches, InverseFaceNet for the first time also includes color reflectance and illumination. We furthermore significantly improve the implicitly learned priors by continuously updating the training corpus based on the distribution of real-world imagery using *breeder learning* inspired by Nair et al. [32]. The key idea is to apply the latest version of the inverse face rendering network to real photographs, to estimate the corresponding face model parameters, and then to create synthetic face renderings for perturbed (but still known) parameter values. This way, we create additional synthetic training data that better reflects the real-world distribution of face model parameters, and our network learns this distribution over time. Our experiments demonstrate that our approach greatly improves the quality of regressed face models for real face images compared to approaches that are trained exclusively on synthetic data.

The main contribution of our paper is InverseFaceNet – a fast, deep, single-shot inverse face rendering network that estimates pose, shape, expression, color reflectance and illumination from just a single input image in a single forward pass, and is multiple orders of magnitude faster than previous optimization-based methods estimating similar models. To improve the accuracy of the results, we further propose a loss function that measures model-space distances directly in a modified parameter space. We further propose analysis-by-synthesis breeder learning to produce labeled training data following real-world parameter distributions, which leads to significantly improved reconstruction results for in-the-wild face photos.

2. Related Work

Reconstruction of faces and their appearance has been studied extensively, so we only review the most relevant subset of this vast literature.

Inverse Rendering (of Faces) The goal of inverse rendering is to invert the graphics pipeline, i.e., to recover the geometry, reflectance (albedo) and illumination from images or videos of a scene – or, in our case, a face. Early work on inverse rendering made restrictive assumptions like known scene geometry and calibrated input images [34, 49]. However, recent work has started to relax these assumptions for specific classes of objects such as faces. Deep neural networks have been shown to be able to invert simple graphics pipelines [24, 32], although these techniques are so far only applicable to low-resolution grayscale images. In contrast, our approach reconstructs full-color facial reflectance and illumination, and geometry. Aldrian and Smith [2] use a 3D morphable model for optimization-based inverse rendering. They sequentially solve for geometry, reflectance and illumination, while we jointly regress all dimensions. Thies et al. [45] recently proposed a real-time inverse rendering approach for faces that estimates a person's identity and expression using a blendshape model with reflectance texture and colored spherical harmonics illumination. Their approach is designed for reenactment and is visually convincing, but relies on non-linear least-squares optimization, which requires good initialization and a face model calibration step from multiple frames, while our approach estimates a very similar face model in a single shot, from a single in-the-wild image, in a fraction of the time. Inverse rendering has also been applied to face image editing [30], for example to apply makeup [26, 27]. However, these approaches perform an image-based intrinsic decomposition without an explicit 3D face model, as in our case.

Face Models The appearance and geometry of faces are often modeled using 3D morphable models [4] or active appearance models [11]. These seminal face models are powerful and expressive, and remain useful for many applications even though more complex and accurate appearance models exist [22]. For example, Booth et al. [5] extend 3D morphable models to "in-the-wild" conditions, and deep appearance models [12] extend active appearance models by capturing geometry and appearance of faces more accurately under large unseen variations. We describe the face model we use in Section 4.

3D Face Reconstruction The literature on reconstructing face geometry, often with appearance, but without any illumination, is much more extensive compared to inverse rendering. We focus on single-view techniques and do not further discuss multi-view or multi-image approaches [17, 37, 43]. Recent techniques approach monocular face reconstruction by fitting active appearance models [1, 12], blendshape models [6, 13, 14], multi-linear face models [35, 36, 41, 46], or mesh geometry [19, 25, 37] to input images or videos. Shading-based surface refinement can extract even fine-scale geometric surface detail [8, 14, 19, 36, 37]. Many techniques use facial landmark detectors for more robustness to changes in the head pose and expression, and we discuss them in the next section. A range of approaches use RGB-D input [e.g. 28, 44, 48], and while they achieve impressive face reconstruction results, they rely on depth data which is typically not available for in-the-wild images or videos.

Deep neural networks have recently shown promising results on various face reconstruction tasks. In a paper before its time, Nair et al. [32] proposed an analysis-by-synthesis algorithm that iteratively explores the parameter space of a black-box generative model, such as active appearance models (AAM) [11], to learn how to invert it, e.g., to convert a photo of a face into an AAM parameter vector. We are inspired by this *breeding* approach and incorporate a similar strategy into our training process (see Section 7) to make our technique more robust to unseen inputs, in our case real photographs. Richardson et al. [35] use iterative error feedback [9] to optimize the shape parameters of a grayscale

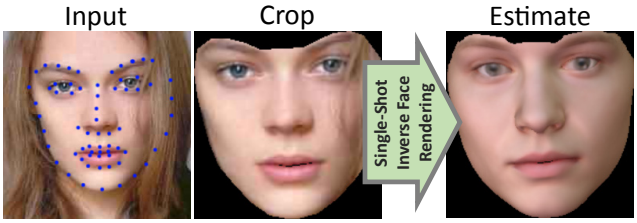

Figure 2. Our single-shot inverse face renderer regresses a dense reconstruction of the pose, shape, expression, skin reflectance and incident illumination from a single input image of a face.

morphable model from a single input image. Richardson et al. [36] build on this to reconstruct detailed depth maps of faces with learned shape-from-shading. Their approaches are trained entirely on synthetic data [4]. Tran et al. [46] train using a photo collection, but their focus lies on estimating morphable model parameters to achieve robust face recognition. In contrast to these approaches, ours not only recovers face geometry and texture, but a more complete inverse rendering model that also comprises color reflectance and illumination, from just a single image without the need for iteration.

Facial Landmark Detection Many techniques in 3D face reconstruction, including ours, draw on facial landmark detectors for robustly identifying the location of landmark keypoints in the photograph of a face, such as the outline of the eyes, nose and lips. These landmarks can provide valuable pose-independent initialization. Chrysos et al. [10] and Jin and Tan [20] provide two recent surveys on the many landmark detection approaches that have been proposed in the literature. However, none of these techniques works perfectly [42]: facial hair, glasses and poor lighting conditions pose the largest problems. In many cases, these problems can be overcome when looking at video sequences instead of single images [33], but this is a different setting to ours.

3. Overview

We propose a single-shot deep inverse face rendering network that learns the pose, shape, expression, skin reflectance and incident illumination from a single input image of a face. To this end, we first detect a set of 66 2D facial landmarks [39], see Figure 2. We use the landmarks to segment the face from the background, and mask out the mouth interior to effectively remove the parts of the image that cannot readily be explained by our model. The masked face is input to our deep inverse face rendering network (Section 6). We generate purely synthetic facial imagery using a parametric face and image formation model (Section 4). Our training corpus is sampled from the probability distribution of this model (Section 5). We apply parameter breeding by analysis-by-synthesis (Section 7) to bootstrap a synthetic training corpus that approximates the real-world distribution of faces and scene illumination. This leads to higher quality reconstruction results (Section 8). Finally, we discuss limitations (Section 8.4) and conclude the paper (Section 9).

4. The Space of Facial Imagery

We parameterize face images using $m = 350$ parameters:

$$\boldsymbol{\theta} = \left(\mathbf{R}, \boldsymbol{\theta}^{[\mathrm{s}]}, \boldsymbol{\theta}^{[\mathrm{e}]}, \boldsymbol{\theta}^{[\mathrm{r}]}, \boldsymbol{\theta}^{[\mathrm{i}]}\right) \in \mathbb{R}^m. \tag{1}$$

Here, $\mathbf{R}$ specifies the global rotation (3 parameters), $\boldsymbol{\theta}^{[\mathrm{s}]}$ the shape (128), $\boldsymbol{\theta}^{[\mathrm{e}]}$ the expression (64), $\boldsymbol{\theta}^{[\mathrm{r}]}$ the skin reflectance (128), and $\boldsymbol{\theta}^{[\mathrm{i}]}$ the incident illumination (27). Note that we do not include translation as our network works on consistently segmented input images (see Figure 2 and Section 3). In the following, we describe all parameters and the underlying models in detail.

4.1. Multi-Linear Face Model

We employ a multi-linear face model to parameterize facial geometry $\mathcal{F}^{[\mathrm{g}]} \in \mathbb{R}^{3V}$ and skin reflectance $\mathcal{F}^{[\mathrm{r}]} \in \mathbb{R}^{3V}$, where V is the number of vertices of the underlying manifold template mesh. The geometry vector $\mathcal{F}^{[\mathrm{g}]}$ stacks the V 3D coordinates that define the mesh's embedding in space. Similarly, the reflectance vector $\mathcal{F}^{[\mathrm{r}]}$ stacks the RGB reflectance values. The space of facial geometry is modeled by the shape $\boldsymbol{\theta}^{[\mathrm{s}]} \in \mathbb{R}^{N_\mathrm{s}}$ and expression $\boldsymbol{\theta}^{[\mathrm{e}]} \in \mathbb{R}^{N_\mathrm{e}}$ parameters:

$$\mathcal{F}^{[\mathrm{g}]}(\boldsymbol{\theta}^{[\mathrm{s}]}, \boldsymbol{\theta}^{[\mathrm{e}]}) = \mathbf{a}^{[\mathrm{g}]} + \sum_{i=1}^{N_\mathrm{s}} \mathbf{b}_i^{[\mathrm{s}]} \sigma_i^{[\mathrm{s}]} \theta_i^{[\mathrm{s}]} + \sum_{j=1}^{N_\mathrm{e}} \mathbf{b}_j^{[\mathrm{e}]} \sigma_j^{[\mathrm{e}]} \theta_j^{[\mathrm{e}]}. \tag{2}$$

The spatial embedding is modeled by a linear combination of orthonormal basis vectors $\mathbf{b}_i^{[\mathrm{s}]}$ and $\mathbf{b}_j^{[\mathrm{e}]}$, which span the shape and expression space, respectively. $\mathbf{a}^{[\mathrm{g}]} \in \mathbb{R}^{3V}$ is the average geometry of a neutral expression, the $\sigma_i^{[\mathrm{s}]}$ are the shape standard deviations and the $\sigma_j^{[\mathrm{e}]}$ are the standard deviations of the expression dimensions.

Per-vertex skin reflectance is modeled similarly using a small number of reflectance parameters $\boldsymbol{\theta}^{[\mathrm{r}]} \in \mathbb{R}^{N_\mathrm{r}}$:

$$\mathcal{F}^{[\mathrm{r}]}(\boldsymbol{\theta}^{[\mathrm{r}]}) \;=\; \mathbf{a}^{[\mathrm{r}]} + \sum_{i=1}^{N_\mathrm{r}} \mathbf{b}_i^{[\mathrm{r}]} \sigma_i^{[\mathrm{r}]} \theta_i^{[\mathrm{r}]}. \tag{3}$$

Here, $\mathbf{b}_i^{[\mathrm{r}]}$ are the reflectance basis vectors, $\mathbf{a}^{[\mathrm{r}]}$ is the average reflectance and the $\sigma_i^{[\mathrm{r}]}$ are the standard deviations.

The face model is computed from 200 high-quality 3D scans [4] of Caucasians (100 male and 100 female) using PCA. We use the $N_\mathrm{s} = N_\mathrm{r} = 128$ most significant principal directions to span our face space. The used expression basis is a combination of the Digital Emily model [3] and FaceWarehouse [7], see Thies et al. [45] for details. We use PCA to compress the over-complete blendshapes (76 vectors) to a subspace of $N_\mathrm{e} = 64$ dimensions.

4.2. Image Formation

We assume the face to be *Lambertian*, illumination to be distant and smoothly varying, and there is no self-shadowing. We thus represent the incident illumination on the face using second-order spherical harmonics (SH) [31, 34]. Therefore, the irradiance at a surface point with normal $\mathbf{n}$ is given by

$$\mathcal{B}\left(\mathbf{n} \mid \boldsymbol{\theta}^{[\mathrm{i}]}\right) = \sum_{k=1}^{b^2} \boldsymbol{\theta}_k^{[\mathrm{i}]} H_k(\mathbf{n}), \tag{4}$$

where H_k are the $b^2 = 3^2 = 9$ SH basis functions, and the $\boldsymbol{\theta}_k^{[\mathrm{i}]}$ are the corresponding illumination coefficients. Since we consider colored illumination, the parameters $\boldsymbol{\theta}_k^{[\mathrm{i}]} \in \mathbb{R}^3$ specify RGB colors, leading to $3 \cdot 9 = 27$ parameters in total.

We render facial images based on the SH illumination using a full perspective camera model $\Pi\colon \mathbb{R}^3 \to \mathbb{R}^2$. We render the face using a mask (painted once in a preprocessing step) that ensures that the rendered facial region matches the crops produced by the 66 detected landmark locations (see Figure 2). The global rotation of the face is modeled with three Euler angles using $\mathbf{R} = \mathrm{Rot}_{xyz}(\alpha, \beta, \gamma)$ that successively rotate around the x-axis (up, α), y-axis (right, β), and z-axis (front, γ) of the camera space coordinate system.

5. Synthetic Training Data Generation

Training our deep inverse face rendering network requires ground-truth training data $\{\mathbf{I}_i, \boldsymbol{\theta}_i\}_{i=1}^N$ in the form of corresponding pairs of image $\mathbf{I}_i$ and model parameters $\boldsymbol{\theta}_i$. However, training on real images is challenging, since the ground-truth parameters cannot easily be obtained for a large dataset. We therefore train our network based on synthetically rendered data, where exact ground-truth labels are available.

We sample $N = 200{,}000$ parameter vectors $\boldsymbol{\theta}_i$ and use the model described in Section 4 to generate the corresponding images $\mathbf{I}_i$. Data generation can be interpreted as sampling from a probability $P(\boldsymbol{\theta})$ that models the distribution of real-world imagery. However, sampling from this distribution is in general difficult and non-trivial. We therefore assume statistical independence between the components of $\boldsymbol{\theta}$, i.e.,

$$P(\boldsymbol{\theta}) = P(\mathbf{R})P(\boldsymbol{\theta}^{[\mathrm{s}]})P(\boldsymbol{\theta}^{[\mathrm{e}]})P(\boldsymbol{\theta}^{[\mathrm{r}]})P(\boldsymbol{\theta}^{[\mathrm{i}]}). \tag{5}$$

This enables us to efficiently generate a parameter vector $\boldsymbol{\theta}$ by independently sampling each subset of parameters.

We uniformly sample the yaw and pitch rotation angles $\alpha, \beta \sim U(-40°, 40°)$ and the roll angle $\gamma \sim U(-15°, 15°)$ to reflect common head rotations in uncontrolled photos. We sample shape and reflectance parameters from the Gaussian distributions provided by the parametric PCA face model [4]. Since we already scale with the appropriate standard deviations during face generation (see Equations 2 and 3), we sample both from a standard normal distribution, i.e., $\boldsymbol{\theta}^{[\mathrm{s}]}, \boldsymbol{\theta}^{[\mathrm{r}]} \sim N(0, 1)$. The expression basis is based on artist-created blendshapes that only approximate the real-world distribution of the space of human expressions; this will be addressed by the breeding step presented in Section 7. We thus uniformly sample the expression parameters using $\boldsymbol{\theta}^{[\mathrm{e}]} \sim U(-12, 12)$. To prevent closing the mouth beyond anatomical limits, we apply a bias of 4.8 to the distribution of the first parameter[1]. Finally, we sample the illumination parameters using $\boldsymbol{\theta}^{[\mathrm{i}]} \sim U(-0.2, 0.2)$, except for the constant coefficient $\boldsymbol{\theta}_1^{[\mathrm{i}]} \sim U(0.6, 1.2)$ to account for the average image brightness, and set all RGB components to the same value to achieve monochrome illumination. The breeding step presented in Section 7 introduces colored illumination.

6. InverseFaceNet

Given the training data $\{\mathbf{I}_i, \boldsymbol{\theta}_i\}_{i=1}^N$ consisting of N images $\mathbf{I}_i$ and the corresponding ground-truth parameters $\boldsymbol{\theta}_i$, we train a deep inverse face rendering network $\mathcal{N}$ to invert image formation. In the following, we provide details on our network architecture and the employed loss function.

6.1. Network Architecture

We have tested several different networks based on the popular AlexNet [23] and ResNet [15] architectures, both pre-trained on ImageNet [38]. In both cases, we resize the last fully-connected layer to match the dimensionality of our model (350 outputs), and initialize biases with 0, and weights $\sim N(0, 0.01)$. These minimally modified networks provide the baseline we build on. We introduce more substantial changes to the training procedure by introducing a novel model-space loss in Section 6.2, which more effectively trains the same network architecture. The color channels of the input images are normalized to the range $[-0.5, 0.5]$ before feeding the data to the network. We show a comparison between the results of AlexNet and ResNet-101 in Section 8.1, and thus choose AlexNet for our results.

Input Pre-Processing The input to our network is a color image of a masked face with a resolution of 240×240 pixels (see Figure 2). We mask the face to remove any background and the mouth interior which cannot be explained by our model. For this, we detect landmarks [39] and resize their bounding box uniformly to fit inside 240×240 pixels, to approximately achieve scale and translation invariance.

Training We train all our inverse face rendering networks using the Caffe deep learning framework [18] with stochastic gradient descent based on AdaDelta [50]. We perform 75K batch iterations with a batch size of 32 for training our baseline approaches. To prevent overfitting, we use an ℓ_2-regularizer (*aka* weight decay) of 0.001. We train with a base learning rate of 0.01.

[1] The first parameter mainly corresponds to mouth opening and closing.

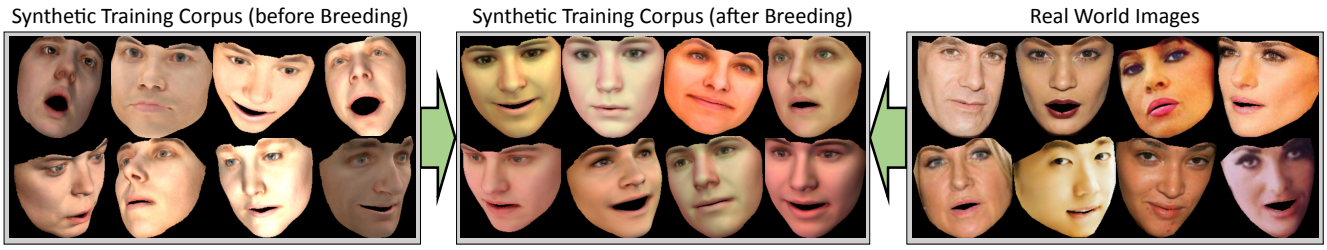

Figure 3. Our approach updates the initial training corpus (left) based on the distribution of real-world imagery (right) using an *analysis-by-synthesis breeding strategy* [32]. The generated new training corpus (middle) better reflects the real-world distribution of facial images.

6.2. Model-Space Parameter Loss

We employ a modified Euclidean loss that measures distance between predicted parameters $\mathbf{p}$ and ground-truth parameters $\boldsymbol{\theta}$ by taking the statistics of the face model into account:

$$\mathcal{L}(\mathbf{p},\boldsymbol{\theta}) = \big\|\mathbf{p}-\boldsymbol{\theta}\big\|^2_{\mathbf{A}} \tag{6}$$

$$= (\mathbf{p}-\boldsymbol{\theta})^\top \underbrace{\mathbf{A}}_{\boldsymbol{\Sigma}^\top\boldsymbol{\Sigma}} (\mathbf{p}-\boldsymbol{\theta}). \tag{7}$$

Here, $\boldsymbol{\Sigma}$ is a weight matrix that incorporates the standard deviations $\boldsymbol{\sigma}^\bullet$ of the different parameter dimensions:

$$\boldsymbol{\Sigma} = \mathrm{diag}(\omega_{\mathbf{R}}\mathbf{1}_3, \omega_{\mathrm{s}}\boldsymbol{\sigma}^{[\mathrm{s}]}, \omega_{\mathrm{e}}\boldsymbol{\sigma}^{[\mathrm{e}]}, \omega_{\mathrm{r}}\boldsymbol{\sigma}^{[\mathrm{r}]}, \omega_{\mathrm{i}}\mathbf{1}_{27}) \in \mathbb{R}^{m\times m}. \tag{8}$$

The coefficients $\omega_\bullet$ balance the global importance of the different groups of parameters, and $\mathbf{1}_k$ is a k-dimensional vector of ones. We use the same values $(\omega_{\mathbf{R}}, \omega_{\mathrm{s}}, \omega_{\mathrm{e}}, \omega_{\mathrm{r}}, \omega_{\mathrm{i}}) = (400, 50, 50, 100, 20)$ for all our results. Note that we do not scale the rotation and illumination dimensions individually, which allows them to vary freely. Intuitively speaking, our model-space loss enforces that the first PCA coefficients (higher variation basis vectors) should match the ground truth more accurately than the later coefficients (lower variation basis vectors), since the former have a larger contribution to the final geometry and skin reflectance of the reconstructed face in model space (see Equations 2 and 3). As shown in Section 8, this leads to more natural and accurate reconstruction results.

7. Breeding with Real Images

The real-world distribution of the model parameters $\boldsymbol{\theta}$ is in general unknown for in-the-wild images. Until now, we have sampled from a manually prescribed probability distribution $P(\boldsymbol{\theta})$, which does not exactly represent the real-world distribution. The goal of our parameter 'breeding' step is to make the training data distribution better match the real-world distribution in a corpus $\mathcal{R}$ of in-the-wild face photographs. To this end, the breeding step automatically bootstraps the parameters for the training corpus $\mathcal{R}$ using analysis-by-synthesis [32]. Note that this step is unsupervised and does not require the ground-truth parameters for images in $\mathcal{R}$ to be available.

Algorithm 1 Breeding with Real Images

$\mathcal{N} \leftarrow$ train_network_on_synthetic_faces();
$\mathcal{R} \leftarrow$ corpus_of_real_images();
for (number of breeding steps N_{breed}) **do**
  $\boldsymbol{\theta}_r \leftarrow$ inverse_rendering($\mathcal{R}, \mathcal{N}$); ▷ (step 1)
  $\boldsymbol{\theta}'_r \leftarrow$ resample_parameters($\boldsymbol{\theta}_r$); ▷ (step 2)
  $\mathbf{I}_r \leftarrow$ generate_images($\boldsymbol{\theta}'_r$); ▷ (step 3)
  $\mathcal{N} \leftarrow$ continue_training($\mathcal{N}, \{\mathbf{I}_r, \boldsymbol{\theta}'_r\}$); ▷ (step 4)
end for

Breeding is a four-step process (see Algorithm 1). It starts with a deep neural network $\mathcal{N}$ initially trained on a synthetic training corpus (Section 5) for 15K batch iterations. This guarantees a suitable initialization for all weights in the network. Given a set of images from the corpus of real-world images $\mathcal{R}$, we first obtain an estimate of the corresponding model parameters $\boldsymbol{\theta}_r$ using the synthetically trained network (step 1). These reconstructed parameters are used to seed the breeding process. In the next step, we apply small perturbations to the reconstructed parameters (step 2) based on a noise distribution. This generates new data around the seed points in the model parameter space, and allows the network to slowly adapt to the real-world parameter distribution. We use the following noise distributions to resample the pose, shape, expression, reflectance and illumination parameters, generating two perturbed parameter vectors for each reconstructed parameter vector: α, β, γ: $U(-5°, 5°)$, $\boldsymbol{\theta}^{[\mathrm{s}]}$: $N(0, 0.05)$, $\boldsymbol{\theta}^{[\mathrm{r}]}$: $N(0, 0.2)$, $\boldsymbol{\theta}^{[\mathrm{e}]}$: $N(0, 0.1)$, and $\boldsymbol{\theta}^{[\mathrm{i}]}$: $N(0, 0.02)$. In the next processing step (step 3), we generate new synthetic training images $\mathbf{I}_r$ based on the resampled parameters $\boldsymbol{\theta}'_r$. The result is a new synthetic training set $\{\mathbf{I}_r, \boldsymbol{\theta}'_r\}$ that better reflects the real-world distribution of model parameters. Finally, the network $\mathcal{N}$ is fine-tuned for $N_{\text{iter}} = 7.5$K batch iterations on the new training corpus (step 4). In total, we repeat this process for $N_{\text{breed}} = 8$ iterations of these breeding steps. Over the iterations, the data distribution of the training corpus adapts and better reflects the real-world distribution of the provided in-the-wild facial imagery, as illustrated in Figure 3. This leads to better reconstruction results at test time, as shown in Section 8.

Table 1. Quantitative comparison of network architectures, model-space parameter loss and our breeding steps on 5,914 test images from *CelebA* [29]. The best values for each column are highlighted in bold. Training time includes all algorithmic steps except the initial training data generation. Test times are averaged over 5K images. All timings on an Nvidia GeForce GTX Titan. Errors show means and standard deviations. *For breeding, we first train 15K iterations on normal synthetic face images (Section 5), and then breed for a further 60K iterations (see Section 7). InverseFaceNet (AlexNet [23] with model-space loss and breeding) produces the best results in all error categories.

Approach	Training iterations	Training time [h]	Test time [ms / image]	Photometric error [8 bits]	Geometric error [mm]	Intersection over union [%]
AlexNet [23]	75K	4.14	**9.79**	46.26 ±12.42	2.91 ±0.99	90.44 ±3.81
+ model-space loss	75K	4.36	**9.79**	39.71 ± 9.86	2.77 ±1.00	92.51 ±2.59
+ breeding (= **InverseFaceNet**)	75K*	29.40	**9.79**	**34.03 ± 7.56**	**2.11 ±0.84**	**93.96 ±2.08**
ResNet-101 [15] + model-space loss	150K	40.99	134.45	41.23 ±10.58	2.54 ±0.87	92.07 ±2.87

8. Experiments and Results

We evaluate our InverseFaceNet on several publicly available datasets. We start by validating our design choices regarding the network architecture (Section 6), model-space parameter loss (Section 6.2), and breeding of the training corpus (Section 7). We then show qualitative and quantitative results and comparisons on the datasets *CelebA* [29], *LFW* (Labeled Faces in the Wild) [16], *FaceWarehouse* [7], *Thomas* [13], *Volker* [47] and *300-VW* (300 Videos in the Wild) [40].

Error Measures We measure the quality of reconstructed face models using photometric error, geometric error, and overlap of face masks. We compute the **photometric error** using the RMSE of RGB pixel values (within the mask of the input image) between the input image and a rendering of the reconstructed face model. An error of 0 is a perfect color match, and 255 is the difference between black and white (i.e. lower is better). The **geometric error** measures the RMSE in mm between corresponding vertices in our reconstruction and the ground-truth geometry. We quantify the image-space overlap of the estimated face model and the input face image using the **intersection over union** (IOU) of face masks (e.g. see 'contours' in Figure 4). An IOU of 0% means no overlap, and 100% means perfect overlap (i.e. higher is better).

8.1. Evaluation of Design Choices

Table 1 evaluates different design choices on a test dataset of 5,914 images (one shown in Figure 4) from *CelebA* [29] using the error measures described earlier (using our implementation of Garrido et al. [14] as ground-truth geometry, up to blendshape level).

Network Architecture We first compare the results of the AlexNet [23] and ResNet-101 [15] architectures, both with our model-space loss (see Section 6). Reconstructions using ResNet-101 have smaller geometric errors, but worse photometric error and IOU than AlexNet, which is exemplified by Figure 4. ResNet-101 is significantly deeper than AlexNet, so training takes about 10× longer (all timings on an Nvidia GeForce GTX Titan). At test time, we reconstruct a full set of face model parameters in 134 ms using ResNet-101, or 9.79 ms using AlexNet, which is 13× faster. We thus use AlexNet for our inverse face rendering network.

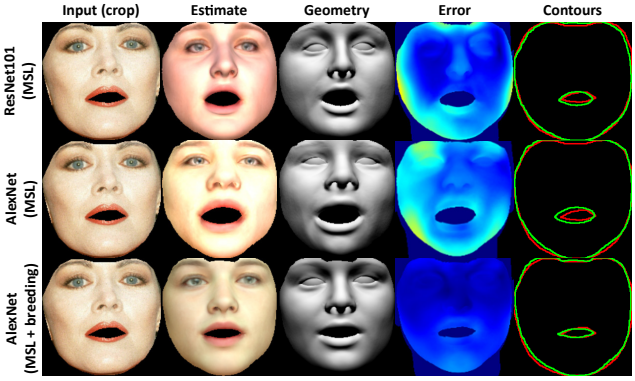

Figure 4. Qualitative comparison of ResNet-101 [15] and AlexNet [23] applied to inverse face rendering, both with model-space loss (MSL): ResNet-101 produces lower geometric error while AlexNet has lower photometric error (also on average, see Table 1). AlexNet with MSL and breeding clearly improves the reconstruction of reflectance and geometry, in all error categories.

Importance of Model-Space Loss Table 1 shows that our model-space loss improves on baseline AlexNet [23] in all error categories, particularly the photometric error and IOU. As our model-space loss does not modify the network architecture, the time to reconstruct a better face model from a single image remains the same fast 9.79 ms as before.

Importance of Breeding Adapting the distribution of face model parameters in the training corpus via breeding (Section 7) significantly improves the reconstruction quality and produces the lowest errors in all categories, as shown in Table 1. This can also be seen in Figure 4, which shows plausible reconstruction of appearance and geometry, the lowest geometric errors, and improved contour overlap for our inverse face rendering network with breeding. Note that the training time for breeding includes all steps of breeding (Algorithm 1), in particular reconstructing 100K face models (0.25 h), rendering 200K synthetic faces (2.8 h) and training for 7.5K iterations (0.5 h) for each of the 8 breeding iterations. AlexNet with breeding significantly outperforms ResNet-101 without breeding in reconstruction quality, train-

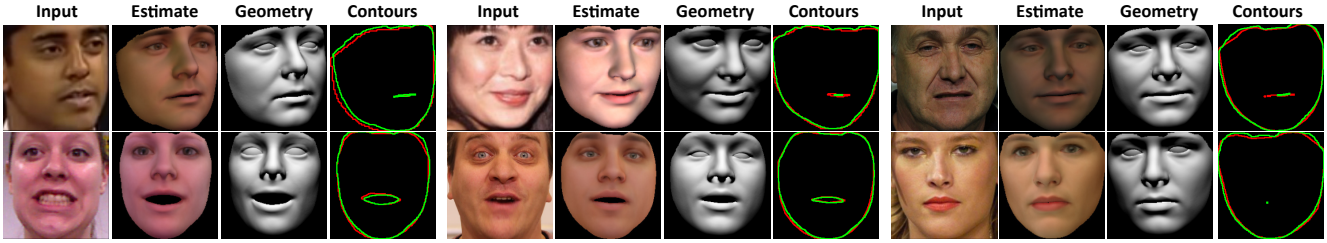

Figure 5. Qualitative results on several datasets. Left to right: input image, our estimated face model and geometry, and contours (red: input mask, green: ours). In reading order: *300-VW* [40], *LFW* [16], *Volker* [47], *FaceWarehouse* [7], *Thomas* [13] and *CelebA* [29]. Our approach achieves high-quality reconstructions of geometry as well as skin reflectance from just a single input image.

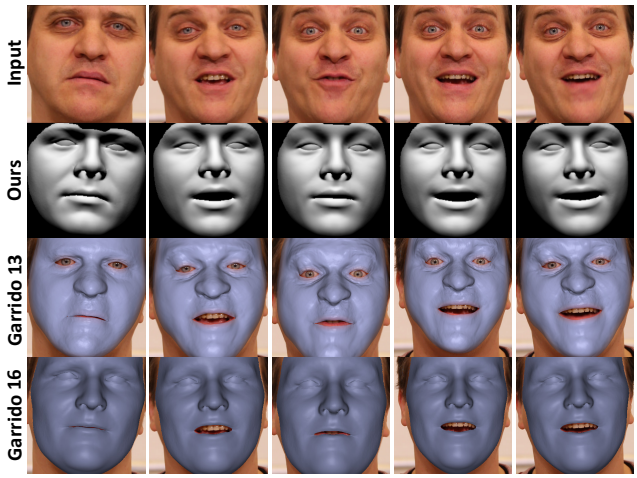

Figure 6. Qualitative comparison to optimization-based approaches [13, 14] on the *Thomas* dataset [13].

ing time and test time, which remains unchanged for AlexNet at 9.79 ms per image, or more than 100 Hz (ResNet: 7.4 Hz).

8.2. Qualitative Evaluation

Figure 5 shows several monocular reconstruction results obtained with our InverseFaceNet. As can be seen, our approach obtains good estimates of all model parameters, i.e., pose, shape, expression, colored reflectance and scene illumination.

We next compare our reconstruction results qualitatively to current state-of-the-art approaches. Figure 6 compares our reconstruction to optimization-based approaches that fit a parametric face model [14] or a person-specific template mesh [13]. Our learning-based approach is significantly faster (9.79 ms vs about one minute [14]), and it is in principle orthogonal to these optimization-based approaches, since it could be used to provide an initial solution close to the optimum.

We also compare to the state-of-the-art deep-learning-based approaches by Richardson et al. [35, 36] and Tran et al. [46]. Specifically, we compare to the results of Richardson et al.'s 'CoarseNet' [36], which corresponds to their earlier method [35], and estimates pose, shape and expression, followed by a model-based optimization of monochrome

Figure 7. Qualitative comparison to Richardson et al. [36] on *LFW* [16]. Note that our reconstruction results are colored, and better fit the face shape and mouth expressions of the input images.

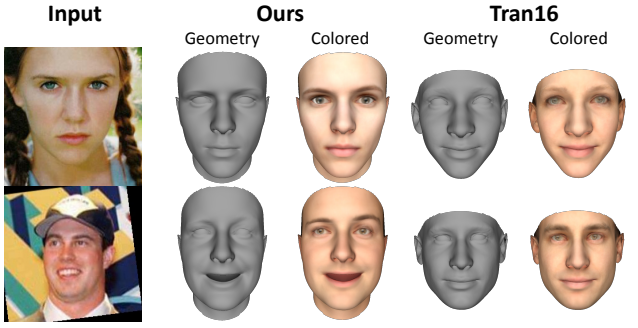

Figure 8. Qualitative comparison to Tran et al. [46] on images of the *CelebA* [29] and *LFW* datasets [16]: our approach reconstructs expressions, while theirs cannot recover this dimension (bottom).

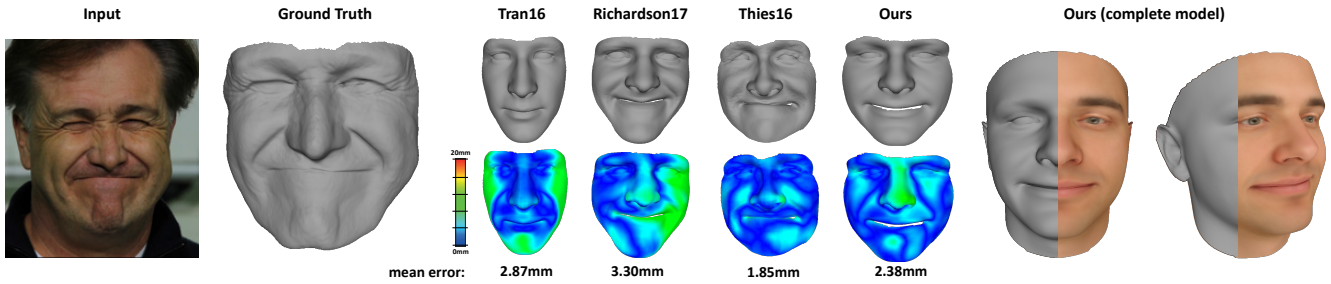

Figure 9. Quantitative comparison of geometric accuracy compared to Tran et al. [46], Richardson et al. [36] and Thies et al. [45] on the *Volker* dataset [47]. The heat maps visualize the pointwise Hausdorff distance (in mm) between the input and the ground-truth mesh. The ground truth has been obtained by the high-quality binocular reconstruction approach of Valgaerts et al. [47].

reflectance and illumination.[2] As can be seen in Figures 7 and 8, we obtain similar or even higher quality results than these two state-of-the-art approaches. In contrast to Richardson et al. [35, 36], our approach also obtains an estimate of the colored skin reflectance and illumination. Note that their approach does not require landmarks for initial cropping, but it is significantly slower due to their iterative regression strategy. The approach of Tran et al. [46] does not recover the facial expression and illumination. In contrast, our approach provides a one-shot estimate of all face model parameters.

8.3. Quantitative Evaluation

We compare the geometric accuracy of our approach to state-of-the-art monocular reconstruction techniques in Figure 9. As ground truth, we use the high-quality stereo reconstructions of Valgaerts et al. [47]. Compared to Thies et al. [45], our approach obtains similar quality results, but without the need for explicit optimization. Therefore, our approach is almost two orders of magnitude faster (9.79 ms vs 600 ms) than optimization-based approaches. Note that while Thies et al. [45] runs in real time for face tracking, it requires significantly longer to estimate all model parameters from an initialization based on the average model. Our reconstructions are also of higher quality than the state-of-the-art learning-based methods by Richardson et al. [35, 36] and Tran et al. [46]. In contrast to these methods, we obtain a reconstruction of more aspects of the same face model. More specifically, we reconstruct pose, shape, expression, and colored skin reflectance and illumination.

8.4. Limitations

We propose a solution to the highly challenging problem of inverse face rendering from a single image. Similar to previous learning-based approaches, ours has a few limitations. Our approach does not perfectly generalize to inputs that are outside of the available training data. This is a limitation our approach shares with many learning-based techniques.

Profile views of the head are problematic and hard to reconstruct, even if they are part of the training corpus. Note that even state-of-the-art landmark trackers often fail in this scenario. Handling these cases robustly remains an open research question. Incorrect landmark localization might produce inconsistent input to our network, which harms the quality of the regressed face model. This could be addressed by either more sophisticated face detection algorithms, or by joint learning of face alignment and reconstruction.

External occlusions of the face, such as hair, beards, sun glasses or hands, can also be problematic. To handle these situations more robustly, our approach could be trained in an occlusion-aware manner by augmenting our training corpus with artificial occlusions, similar to Zhao et al. [51].

Our inverse face rendering network produces temporally quite stable reconstructions if applied to video, as shown in our supplemental video. However, we cannot guarantee temporal coherence on a video sequence as the primary goal of our approach is single-shot monocular reconstruction. This could be alleviated by an additional optimization that ensures temporal coherence, or by training explicitly on temporal data.

9. Conclusion

We have presented InverseFaceNet – a single-shot inverse rendering framework for faces from a single image. Our technique successfully learns the inverse mapping from face images to an elaborate set of face and lighting model parameters that could not be regressed by previous methods. A key contribution of our work is to overcome the lack of well-annotated image datasets by breeding a synthetic training corpus that captures the real-world distribution of facial imagery. This enables high-quality face reconstruction from just a single monocular image. Our evaluation shows that our approach compares favorably to the state of the art, and extends it by also learning colored skin reflectance and scene illumination.

Our learning-based single-shot inverse renderer could be used to quickly and robustly initialize optimization-based monocular reconstruction approaches close to the global minimum. We hope that our approach will stimulate future work in this exciting field.

[2] We don't compare to their 'FineNet', which estimates a refined depth map, as we are interested in comparing to reconstructed face models.

Acknowledgements We thank Anh Tuan Tran and colleagues for making their source code publicly available, and Elad Richardson for running his method on our images. We also thank True-VisionSolutions Pty Ltd for kindly providing the 2D face tracker. This work was supported by the ERC Starting Grant CapReal (335545).

References

[1] J. Alabort-i Medina and S. Zafeiriou. A unified framework for compositional fitting of active appearance models. *IJCV*, 121(1):26–64, 2017. 2

[2] O. Aldrian and W. A. P. Smith. Inverse rendering of faces with a 3D morphable model. *IEEE TPAMI*, 35(5):1080–1093, 2013. 1, 2

[3] O. Alexander, M. Rogers, W. Lambeth, J.-Y. Chiang, W.-C. Ma, C.-C. Wang, and P. Debevec. The Digital Emily Project: Achieving a photorealistic digital actor. *IEEE Computer Graphics and Applications*, 30(4):20–31, 2010. 3

[4] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, pages 187–194, 1999. 2, 3, 4

[5] J. Booth, E. Antonakos, S. Ploumpis, G. Trigeorgis, Y. Panagakis, and S. Zafeiriou. 3D face morphable models "in-the-wild". arXiv:1701.05360, 2017. 2

[6] C. Cao, Y. Weng, S. Lin, and K. Zhou. 3D shape regression for real-time facial animation. *ACM ToG*, 32(4):41:1–10, 2013. 2

[7] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou. FaceWarehouse: A 3D facial expression database for visual computing. *IEEE TVCG*, 20(3):413–425, 2014. 3, 6, 7

[8] C. Cao, D. Bradley, K. Zhou, and T. Beeler. Real-time high-fidelity facial performance capture. *ACM ToG*, 34(4):46:1–9, 2015. 2

[9] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. In *CVPR*, 2016. 2

[10] G. G. Chrysos, E. Antonakos, P. Snape, A. Asthana, and S. Zafeiriou. A comprehensive performance evaluation of deformable face tracking "in-the-wild". arXiv:1603.06015, 2016. 3

[11] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE TPAMI*, 23(6):681–685, 2001. 2

[12] C. N. Duong, K. Luu, K. G. Quach, and T. D. Bui. Deep appearance models: A deep boltzmann machine approach for face modeling. arXiv:1607.06871, 2016. 2

[13] P. Garrido, L. Valgaerts, C. Wu, and C. Theobalt. Reconstructing detailed dynamic face geometry from monocular video. *ACM ToG*, 32(6):158:1–10, 2013. 2, 6, 7

[14] P. Garrido, M. Zollhöfer, D. Casas, L. Valgaerts, K. Varanasi, P. Pérez, and C. Theobalt. Reconstruction of personalized 3D face rigs from monocular video. *ACM ToG*, 35(3):28:1–15, 2016. 1, 2, 6, 7

[15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4, 6

[16] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, 2007. 6, 7

[17] A. E. Ichim, S. Bouaziz, and M. Pauly. Dynamic 3D avatar creation from hand-held video input. *ACM ToG*, 34(4):45:1–14, 2015. 2

[18] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. arXiv:1408.5093, 2014. 4

[19] L. Jiang, J. Zhang, B. Deng, H. Li, and L. Liu. 3D face reconstruction with geometry details from a single image. arXiv:1702.05619, 2017. 2

[20] X. Jin and X. Tan. Face alignment in-the-wild: A survey. arXiv:1608.04188, 2016. 3

[21] I. Kemelmacher-Shlizerman and S. M. Seitz. Face reconstruction in the wild. In *ICCV*, pages 1746–1753, 2011. 1

[22] O. Klehm, F. Rousselle, M. Papas, D. Bradley, C. Hery, B. Bickel, W. Jarosz, and T. Beeler. Recent advances in facial appearance capture. *Computer Graphics Forum*, 34(2):709–733, 2015. 2

[23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012. 4, 6

[24] T. D. Kulkarni, W. Whitney, P. Kohli, and J. B. Tenenbaum. Deep convolutional inverse graphics network. In *NIPS*, 2015. 2

[25] S. Laine, T. Karras, T. Aila, A. Herva, and J. Lehtinen. Facial performance capture with deep neural networks. arXiv:1609.06536, 2016. 2

[26] C. Li, K. Zhou, and S. Lin. Intrinsic face image decomposition with human face priors. In *ECCV*, pages 218–233, 2014. 1, 2

[27] C. Li, K. Zhou, and S. Lin. Simulating makeup through physics-based manipulation of intrinsic image layers. In *CVPR*, pages 4621–4629, 2015. 2

[28] H. Li, J. Yu, Y. Ye, and C. Bregler. Realtime facial animation with on-the-fly correctives. *ACM ToG*, 32(4):42:1–10, 2013. 2

[29] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015. 6, 7

[30] J. Lu, K. Sunkavalli, N. Carr, S. Hadap, and D. Forsyth. A visual representation for editing face images. arXiv:1612.00522, 2016. 2

[31] C. Müller. *Spherical harmonics*. Number 17 in Lecture Notes in Mathematics. Springer, 1966. 4

[32] V. Nair, J. Susskind, and G. E. Hinton. Analysis-by-synthesis by learning to invert generative black boxes. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*, pages 971–981, 2008. 2, 5

[33] X. Peng, R. S. Feris, X. Wang, and D. N. Metaxas. A recurrent encoder-decoder network for sequential face alignment. In *ECCV*, 2016. 3

[34] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH*, pages 117–128, 2001. 1, 2, 4

[35] E. Richardson, M. Sela, and R. Kimmel. 3D face reconstruction by learning from synthetic data. In *3DV*, 2016. 1, 2, 7, 8

[36] E. Richardson, M. Sela, R. Or-El, and R. Kimmel. Learning detailed face reconstruction from a single image. In *CVPR*, 2017. arXiv:1611.05053. 1, 2, 3, 7, 8

[37] J. Roth, Y. Tong, and X. Liu. Unconstrained 3D face reconstruction. In *CVPR*, pages 2606–2615, 2015. 2

[38] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 4

[39] J. M. Saragih, S. Lucey, and J. F. Cohn. Deformable model fitting by regularized landmark mean-shift. *IJCV*, 91(2):200–215, 2011. 3, 4

[40] J. Shen, S. Zafeiriou, G. G. Chrysos, J. Kossaifi, G. Tzimiropoulos, and M. Pantic. The first facial landmark tracking in-the-wild challenge: Benchmark and results. In *ICCV Workshops*, 2015. 6, 7

[41] F. Shi, H.-T. Wu, X. Tong, and J. Chai. Automatic acquisition of high-fidelity facial performances using monocular videos. *ACM ToG*, 33(6):222:1–13, 2014. 2

[42] A. Steger, R. Timofte, and L. V. Gool. Failure detection for facial landmark detectors. arXiv:1608.06451, 2016. 3

[43] S. Suwajanakorn, I. Kemelmacher-Shlizerman, and S. Seitz. Total moving face reconstruction. In *ECCV*, 2014. 2

[44] J. Thies, M. Zollhöfer, M. Nießner, L. Valgaerts, M. Stamminger, and C. Theobalt. Real-time expression transfer for facial reenactment. *ACM ToG*, 34(6):183:1–14, 2015. 2

[45] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Face2Face: Real-time face capture and reenactment of RGB videos. In *CVPR*, 2016. 1, 2, 3, 8

[46] A. T. Tran, T. Hassner, I. Masi, and G. Medioni. Regressing robust and discriminative 3D morphable models with a very deep neural network. arXiv:1612.04904, 2016. 2, 3, 7, 8

[47] L. Valgaerts, C. Wu, A. Bruhn, H.-P. Seidel, and C. Theobalt. Lightweight binocular facial performance capture under uncontrolled lighting. *ACM ToG*, 31(6):187:1–11, 2012. 6, 7, 8

[48] T. Weise, S. Bouaziz, H. Li, and M. Pauly. Realtime performance-based facial animation. *ACM ToG*, 30(4):77:1–10, 2011. 2

[49] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *SIGGRAPH*, pages 215–224, 1999. 2

[50] M. D. Zeiler. ADADELTA: An adaptive learning rate method. arXiv:1212.5701, 2012. 4

[51] F. Zhao, J. Feng, J. Zhao, W. Yang, and S. Yan. Robust LSTM-autoencoders for face de-occlusion in the wild. arXiv:1612.08534, 2016. 8